Fermi National Accelerator Laboratory

TM-1489

Updated Overview of the TEVATRON Control System*

P. Lucas
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

October 1987

*Submitted to the Europhysics Conference on Control Systems for Experimental Physics at CERN, Geneva, Switzerland, September 28-October 2, 1987

Operated by Universities Research Association Inc. under contract with the United States Department of Energy

UPDATED OVERVIEW OF THE TEVATRON CONTROL SYSTEM*

P. Lucas
Fermilab MS307, P.O. Box 500
Batavia, IL 60510 USA

Abstract: A single unified control system is used for all of the Fermilab accelerators and storage rings, from the LINAC to the Tevatron and antiproton source. A review of the general features is given - these include a 'host' system consisting of a number of minicomputers integrated with many distributed microprocessors in a variety of subsystems, usage of an in-house developed protocol, GAS, for communication between the two classes of machines, and a Parameter Page program, designed in conjunction with the system database, which allows a wide variety of quantities to be read and set in a coherent fashion. Recent developments include the implementation of a block transfer and 'fast time plot' facility through CAMAC, inclusion of several new computers in the host, a better understanding of system throughput, greatly improved reliability, advent of programs which sequence a large number of independent operations, and the construction of new hardware subsystems. Possible future system upgrades will be briefly presented. A summary of the utilization of a quite large software staff, at a time when the system is no longer under construction, will be discussed in an appendix as a topic of interest to many attendees of this conference.

Introduction

As part of the construction of the superconducting Tevatron all the Fermilab accelerators, from the LINAC pre-accelerator to the antiproton accumulator and fixed target switchyard, were put under a unified control system, ACNET. This system, as it existed in 1985 in its nearly completed form, has been described in considerable detail elsewhere[1]. The purpose of this paper is to provide an extremely brief summary of that earlier one, to present in detail the upgrades performed from 1985 to 1987, to summarize the system utilization in terms of network transfers and VAX computer cycles , and to indicate directions which future upgrades are likely to take. An appendix is included detailing the current activities of software personnel.

Brief System Description

A block diagram of the Tevatron control system is presented as Figure 1. Much of the figure contains VAX and PDP-11 minicomputers, and the interconnections among them. These computers are collectively referred to as the the 'host system'. The PDP-11's are used as console engines, one per console, and as drivers for the CAMAC, Ethernet, Token Ring, and PDC links to hardware and microprocessors. Two of the VAXes connected as a cluster serve as the network central node - storing all the programs run anywhere on the system, as well as the central database and data files used by various applications. Other functions are to run applications such as data logging which function continuously, and to switch messages among the three separate PCL (Parallel Communication) links networking the computers.

The actual control hardware consists of modules containing varying amounts of intelligence. There are boards, such as digital status and control modules, which perform their functions without the aid of microprocessors; there are modules, such as timing system interfaces and A/D controllers, which contain processors but communicate through standard CAMAC commands; and there are more intelligent processor-based systems primarily housed in Multibus. These 'smart' subsystems perform sophisticated functions such as refrigeration control and beam position monitor setup and data reduction. The intelligent systems communicate with the host over the standard hardware links, and use an in-house developed protocol called GAS.

All addressing and other information needed to allow front-ends to access given devices is contained in the central system database. (In general stored data from modules are kept in files controlled by application programs. These files are separate from the database.) One particular application, called the Parameter Page, was designed in conjunction with the database so that it is able to display readings, settings, and status, and send out settings and control bits, for any of about 40 000 scalar devices in the system - all through a single straightforward user interface. A plot package, allowing any system quantity to be plotted as a function either of time or of another quantity, is also standard.

Precise hardware synchronisation - needed for beam transfers and RF manipulations - is achieved via three high frequency clock systems on which appropriate 'events' can be encoded. One of these systems, TCLK, has a 10 MHz asynchronous carrier; the other two, MRBS and TVBS, have carriers synchronous with Main Ring and Tevatron RF systems, and thus beam.

Recent Additions and Improvements

Block Transfer Return

The ability to use block transfers over the CAMAC links was part of the original specification of the control system, and has been completed recently to great advantage. Most of the functionality resides in the CAMAC serial crate controller which utilizes a separate data path from that of the normal single word transfers. The feature functions as follows: when the front end computer is requested to acquire data, it makes a determination based on the number of words to be transferred and the nature of the source module, as to whether a scalar or block transfer method will be faster. If block mode is chosen, the appropriate commands are sent to the crate controller, which proceeds to move the data from module to link to front end. The hardware operates at $.25 \times 10^6$ (3-byte) CAMAC words per second, and can transfer vectors of up to 64K such words. The most significant application of this access method is in the making of large numbers of fast time plots. For high plot rates, up to 720 Hz, the appropriate data are collected by a microprocessor-based CAMAC 190 module, which serves as a buffer. All the data collected at high rates, 15Hz and above, are returned to the host in large blocks, at 3Hz. Although the fast plotting and block transfer capabilities are independent, the former would easily saturate the system if not used in conjunction with the latter.

*Operated by Universities Research Association under contract with the U.S. Department of Energy.

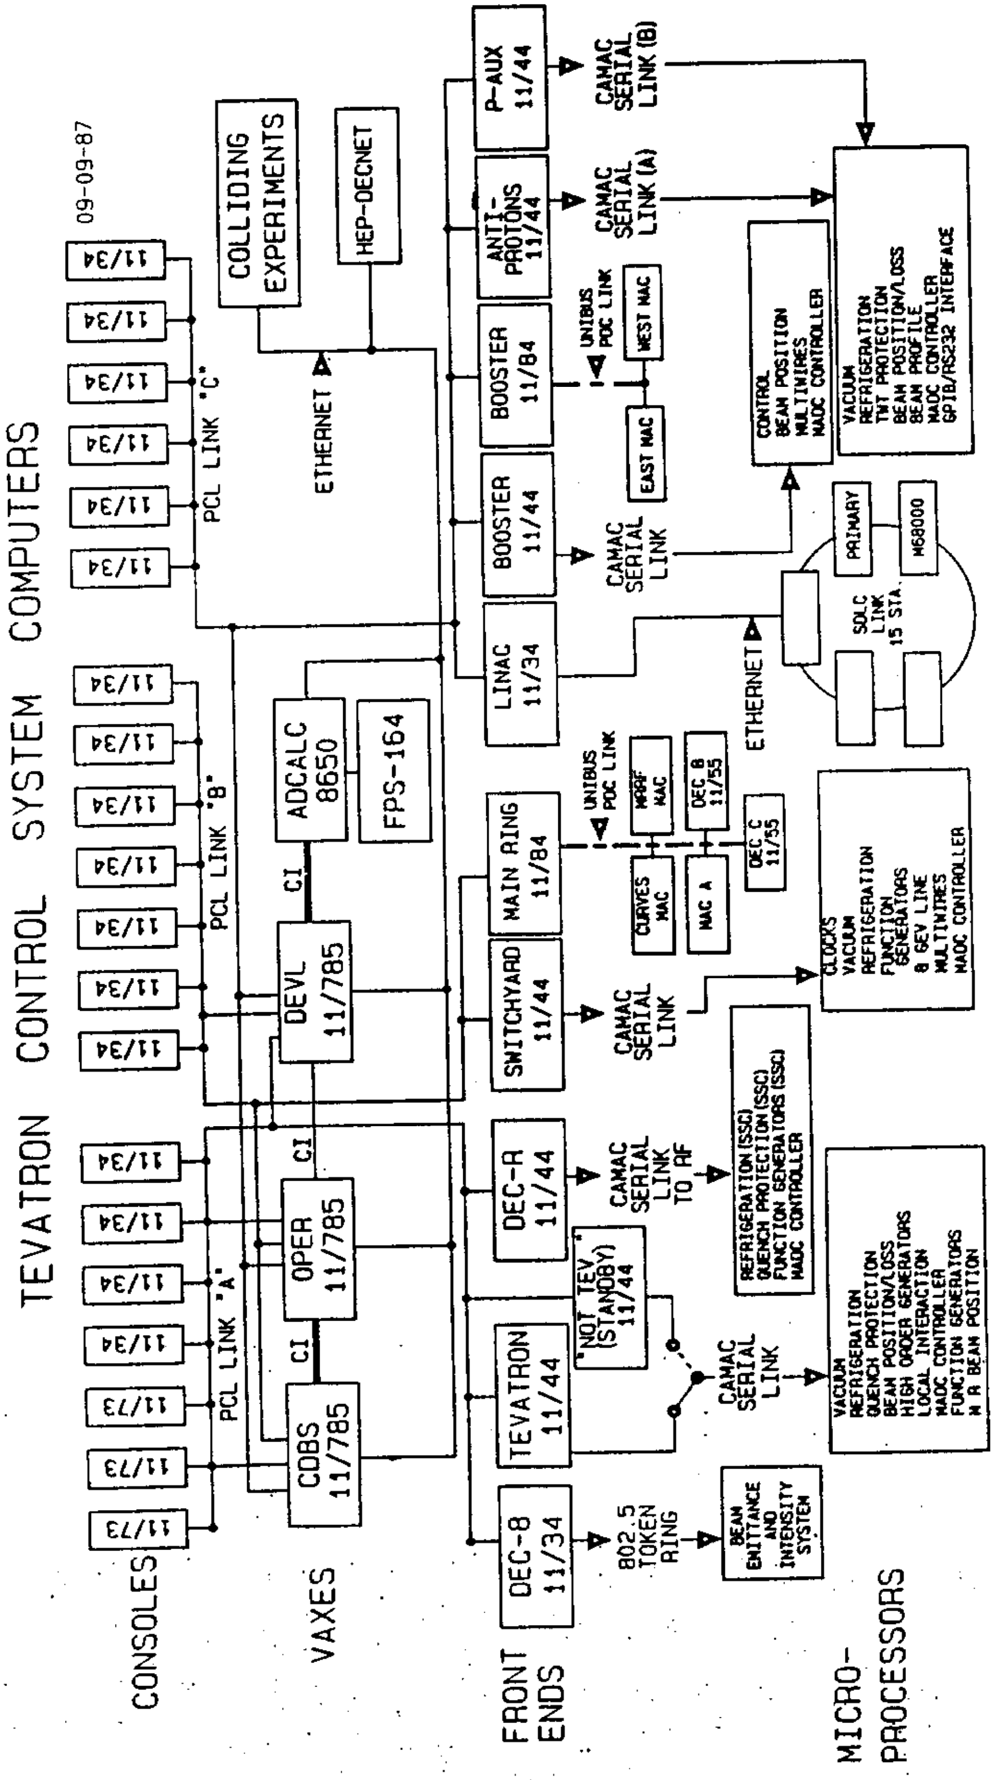
TEVATRON CONTROL SYSTEM COMPUTERS
09-09-87
CONSOLES
VAXES
FRONT ENDS
MICRO-PROCESSORS
11/73
11/73
11/73
11/34
PCL LINK "A"
PCL LINK "B"
PCL LINK "C"
CDBS 11/785
OPER 11/785
DEVL 11/785
CI
CI
CI
ADCALC 8650
FPS-164
COLLIDING EXPERIMENTS
HEP-DECNET
ETHERNET
DEC-8 11/34
TEVATRON 11/44
"NOT TEV" (STANDBY) 11/44
DEC-R 11/44
SWITCHYARD 11/44
MAIN RING 11/84
LINAC 11/34
BOOSTER 11/44
BOOSTER 11/84
ANTI-PROTONS 11/44
P-AUX 11/44
802.5 TOKEN RING
BEAM EMITTANCE AND INTENSITY SYSTEM
CAMAC SERIAL LINK
VACUUM REFRIGERATION QUENCH PROTECTION BEAM POSITION/LOSS HIGH ORDER GENERATORS LOCAL INTERACTION MADC CONTROLLER FUNCTION GENERATORS M R BEAM POSITION
CAMAC SERIAL LINK TO RF
REFRIGERATION (SSC) QUENCH PROTECTION (SSC) FUNCTION GENERATORS (SSC) MADC CONTROLLER
CAMAC SERIAL LINK
CLOCKS VACUUM REFRIGERATION FUNCTION GENERATORS 8 GEV LINE MULTIWIRES MADC CONTROLLER
UNIBUS PDC LINK
CURVES MAC
MRRF MAC
MAC A
DEC B 11/55
DEC C 11/55
ETHERNET
SDLC LINK 15 STA.
PRIMARY
M68000
CAMAC SERIAL LINK
CONTROL BEAM POSITION MULTIWIRES MADC CONTROLLER
EAST MAC
WEST MAC
UNIBUS PDC LINK
CAMAC SERIAL LINK (A)
CAMAC SERIAL LINK (B)
VACUUM REFRIGERATION TWT PROTECTION BEAM POSITION/LOSS BEAM PROFILE MADC CONTROLLER GPIB/RS232 INTERFACE

Figure 1

Additional Computers in the Host

DEC-R: One of the most serious bottlenecks in the system is DEC-T, the Tevatron front end. Various schemes have been examined in an attempt to reduce the load on this machine by splitting it among several engines. The split which seems to have the fewest negative repercussions is to give to all modules residing in or near the RF building a separate front end, since named DEC-R. This split is being implemented at present and is expected to reduce DEC-T's load by 25 - 30%.

DEC-8: The purpose of this front end is to drive an IEEE 802.5 Token Ring link. The first utilization, which will serve as a prototype for future microcomputer devices, is a new beam intensity monitor system. This is being constructed with Token Ring, rather than CAMAC, communication links so that relevant information can be passed between separate processors without going through the host system. This is necessary in particular to allow accurate intensity information to be collected and monitored even during periods in which the host, for whatever reason, is not functional. DEC-8, as indicated in Figure 1, is the interface to ACNET from the Token Ring.

P-AUX: This front end computer is an auxiliary driver of the CAMAC link for the antiproton source. Its implementation is part of a project of converting the Accumulator ring from fixed energy (8 GeV) to variable as part of a new physics program. To effect the necessary energy changes (decelerations) it is necessary to ramp about 100 power supplies which formerly had operated DC. Standard ramp generator cards, as used in the other Fermilab accelerators, are quite expensive and operate much faster than is needed for the majority of devices in the Accumulator. Furthermore the connection of so many power supplies to new control hardware is a large undertaking. Thus it was decided to develop the ramps for such devices inside a single computer and transmit the settings to them over the CAMAC link. The normal antiproton front end DEC-P is inappropriate, since it runs many tasks and has unpredictable real time performance. Thus a new one, P-AUX, dedicated to the ramping task, has been installed. The possibility of having two computers communicating with the same CAMAC devices exists due to the dual port nature of the crate controllers.

CDBS VAX: The functionality of serving as the Accelerator Division Calculational VAX (ADCALC) has been taken over by a new VAX-8650 more suitable for the task than the VAX-11/785 previously used. This latter computer is now clustered with the Operational machine to form a more powerful central node than could be obtained with a single VAX. The new computer is called CDBS, which stands for Central DataBase Services. The original intention in upgrading to a two-machine cluster was that all features of the Operational system involving the database would be split off onto this new engine. In actual implementation of the cluster tasks have been moved between the two machines, indeed onto the Development VAX as well (see Figure 1) at times, to balance the load and maximize throughput. Before this split of the Operational system, a serious saturation condition was often being encountered; at the present time this is not a problem at the central node.

Console engines: Three new consoles have been added since the previous report; the total now stands at twenty.

New Software

Sequencer programs: The most discussed, and most complicated, program of this type is that which orchestrates a 'shot' of antiprotons from the Accumulator back into the Main Ring and Tevatron[2]. This sequenced operation assures that the large number of steps necessary for such an involved procedure are performed successfully and in the proper order. Similar projects on a smaller scale are the automatic start-up of the Main Ring and Tevatron ramps, projects which have been completed, and the automatic pump-down of the Tevatron vacuum, which project is under consideration.

A point being discussed is whether a modern computing technique, such as utilization of a rule-based system, would be appropriate for writing and maintaining sequencers.

MINERVA: This is a user-friendly database entry and editing program. The central database of ACNET was developed in-house and does not use any commercially available database products. One problem with this approach has been that until recently data were entered and modified utilizing a system arcane enough to require experts. This situation contrasts poorly with modern commercial products where truly minimal expertise is required. The data entry program MINERVA represents a step in bringing our database up to current industry standards.

ACNET-60: This process is an extension of normal ACNET transmissions to allow accelerator operating parameters to be made available to experimenters' computers. With the Tevatron operating in collider, as opposed to fixed-target, mode it is important that a number of such parameters be made available to experimenters in a form appropriate to be included on data tapes. These parameters are needed so that luminosity and detailed beam conditions can be determined; in particular the latter are necessary in a proper assesment of detector performance.

The ACNET-60 service establishes cooperating processes in the the experimenters' and the ACNET host system VAXes, the processes communicating via DECNET. The host system VAX process organizes the experimenters' requests and presents them to ACNET in a standard form. The returned data are then distributed to the requesting experiments. This system has given experimenters complete generality in submitting request lists for data.

System reliability improvements: Crashes, usually associated with an abort of one or more processes in the Operational VAX, were a problem with ACNET from its inception. As the system grew and heavier demands were placed upon it, such crashes became both more frequent and less tolerable. The splitting of the functions of that VAX to avoid saturation, and hard work in isolating subtle errors in VAX code, have eliminated several sources of problems. Mean time to failure is now in the range of several days.

New Hardware

CAMAC 365 card: This card is designed to encompass and expand upon the functionality of all previous ramp generators in use at Fermilab. It uses a 10 MHz Z8002 processor to develop a curve of the form,

$$V_{out} = C_1 \cdot g(M_a) + C_2 \cdot h(M_b) + C_3 \cdot f(t)$$

where M_a and M_b are parameters (such as dipole excitation current and slope thereof) and t is

Table 1 - Statistics

a. CAMAC links

Link	Normal CAMAC operations/sec	Block Transfer words/sec	Total bytes/sec	Fraction of maximum
Booster	257	15	544	23%
Tevatron	565	808	2740	41%
Switchyard	775	361	2272	70%

b. VAX computers

VAX	Cycles utilized		I/O operations/sec		Active processes	
	average	max	average	max	average	max
OPER	32%	80%	41	90	33	36
CDBS	15%	90%	11	71	22	24
DEVL	60%	100%	106	281	30	48
ADCALC	30%	94%	30	151	36	68

time. The module allows 15 tables for each of the three functions g, h, and f, and similarly large numbers of choices for scale factors C_1, C_2, and C_3. Function tables and scale factors are selected in various combinations based on events of the real-time clock. Null tables can be selected naturally so that a less general functionality is easily obtained.

This module was originally specified to satisfy the rerquirements of a new low level RF system for the Main Ring. However it is expected to be useful for most new ramp generation applications.

Motion Limit Monitor/Control, MLM/C, system: As noted above, the advent of colliding beam physics in the Tevatron has introduced a much closer coupling of experiments to accelerator operation than existed in the past. There currently are two places where moveable experimental equipment is installed inside the Tevatron vacuum chamber, the purpose of which is to study scattering at angles as small as possible by positioning detectors quite close to the beam. The Accelerator Division has the responsibility of protecting the beam, particularly the antiprotons, from accidental scraping by such equipment. Similarly it is necessary to assure that such devices, while stationary, do not become aperture restrictions when beam conditions change. These goals should be met while interfering as little as possible with the legitimate utilization of the equipment. The MLM/C, a microcomputer (Motorola 68000) based Multibus configured GAS device, has satisfied these two often conflicting goals. Basically, at the beginning of a run the motion of the apparatus is placed in the hands of the accelerator operators, who determine a safe position range for the given beam conditions. Control is then given to the experimenters with the module serving as a watchdog, not allowing any requests for positioning beyond the safe region to be acted upon.

System Utilisation Statistics

A recent accomplishment of some importance has been the ability to monitor in a detailed fashion the amounts of data transmitted over various network links and to correlate such data with those on computer usage. Such a large system as ACNET can only be diagnosed and understood with the help of a rather large statistical base. Note in particular that it is stressed to the maximum during the minutes surrounding an antiproton shot, and must be capable of meeting the demands placed upon it at such a time. Already some bottlenecks and networking errors have been discovered through these studies, with the bottlenecks alleviated and the errors corrected. As the system continues to expand, it is hoped that future problems will be anticipated before they become impediments.

Some statistics on typical system operation are presented here. These values should, however, be interpreted as order of magnitude quantities only as there are several variables which can affect system performance and throughput in a significant fashion. Among these are collider vs fixed target running, time of day, study period vs normal operations, and, during collider operation, shot time vs store time. There is also a difference between instantaneous and time averaged quantities.

CAMAC link utilization as measured by front end computers

Data are presented in Table 1a under the conditions of a long time average during normal fixed target operation. Note that meaningful values cannot be presented for the antiproton front end in this case and also that the installation of DEC-R (see above) was not complete at the time these data were taken. The maximum transmission rates on these links are governed by the computers which drive them; only a small fraction of the link hardware bandwidth itself is ever utilized.

No data are presented on the non-CAMAC front ends. These machines operate in a pool mode in which essentially all available data are collected continuously at a fixed rate. Thus their links are not subject to much variability, and once operating successfully are not modified. The CAMAC links contain the connections to modern microprocessor systems and are thus considerably more complex.

PCL link throughput

During fixed target operation the mean packet transfer rate is 139/sec, the packets varying in size from a minimum of about ten bytes to a maximum of a few hundred. The maximum rate ever observed during fixed target operation is 550/sec averaged over a ten minute period. The comparable record in collider mode is 736/sec. It should be noted that many data transfers in the computer system are synchronous with the accelerator real time clock so that the instantaneous transfer rates immediately following certain clock events considerably exceed the time averages.

VAX computer utilisation

Representative statistics averaged over one day during fixed target operation, and the maximum values for a twelve minute period of that day, are given in Table 1b. The demands on the central node - OPER/CDBS - are observed to be higher on average during collider than fixed target operation, but indications are that there are still considerable spare cycles. Note also that the ADCALC machine is sometimes used for long compute intensive simulation runs, during which all cycles are utilized for considerable periods.

Upgrade Scenarios

As the data processing industry continues to evolve it is both wise and necessary to continually examine systems such as ACNET with a goal of upgrading to modern equipment. This is wise in that modern software packages and hardware can offer improved functionality while requiring minimal manpower. Similarly such upgrades can become necessary to avoid hardware obsolescence. Particular problems faced currently in this regard are that neither MAC computers nor certain console display and hardcopy units are any longer available for purchase, and are repaired only with difficulty. There are also problems, though not too severe yet, with PCL links in the host system. While PCL hardware may still be purchased, and is maintained acceptably, this technology is by and large unsupported on new computers. An example of an improvement which would be wise, although probably not necessary at the present time, is the phasing out of the PDP-11's in the system. Modern equipment, for instance the Microvax line, would provide considerable improvement at minimal cost and eliminate the serious problems associated with the small PDP-11 address space. (Note however that Microvaxes cannot be connected via PCL links.)

Table 2 - Ideas for Next Generation

Consoles:
- Single screen windowing system
- Must run existing multiscreen programs
- Goal of adding 50 new units

Possible solution - Microvax workstation with x-windows under VMS

Front ends:
- Current system bottleneck, likely to become more so
- Too much PDP-11 code to start over

Possible solution - PDP-11/73 chip residing in VME bus. Could distribute computing and multiplex CAMAC links

Network:
- PCL must go, not supported on modern processors
- Network of computers in one room, network of microprocessors in the field. What should be the connection?

Current discussion - IEEE 802.5 Token Ring for microprocessors. DEC suggests Ethernet for computers, we worry about synchronous data and collisions

Distributed systems:
- Need standard languages and uniform communications
- It should be easy from the network point of view to add new types of microprocessors or subsystems
- Can/should one be headed toward Object Oriented Communications?

It is with such considerations in mind that four working groups have been formed in the Controls Department to study the current situation and plan future system upgrades. These four groups are studying the topics Consoles, Front Ends, Networks, and Distributed Systems. A brief summary of their deliberations is presented in Table 2. The possible solutions mentioned in that table represent the thinking as of the present time (September 17, 1987); vigorous discussion continues.

It should be noted that there is a strong constraint on any upgrade of the Fermilab accelerator control system - that it not cause any serious disruption of operations. It is intended that there be no major shutdowns of the facility for the next several years. Thus any changeover step in a Controls upgrade process must be accomplished in at most a few weeks.

Appendix - Software Staff Utilization

The software staff of the Accelerator Controls Department consists of 24 persons, most of them full time data processing professionals. Given a staff of this size some remarks are in order concerning the software effort necessary to create the Tevatron control system, and the effort necessary to maintain it. During the construction phase the Controls software staff provided less than one half of the total effort, the rest coming from elsewhere in the laboratory. However in the present maintenance and upgrade phase most of the effort is indeed supplied by Controls. The current status is that there are so many requests for software that the backlog is over one year, even if no time were spent in planning major system upgrades. Presented in Table A1 is a summary of requests outstanding at the present time. With the exception of those projects associated with the DO experiment, which has been administratively placed in the Fermilab Accelerator Division, this table represents a lengthy list of software upgrades on projects which have been traditionally at Fermilab been considered part of AD/Controls. system

Table A1 - Current Software Projects

There are currently 88 outstanding projects, with breakdown:

Controls network management	9
Application code creation	10
Application code upgrades	23
GPIB interfaces	7
Microprocessor code creation	4
Microprocessor code upgrades	7
Front end upgrades	17
Laboratory fire system	2
Proton, neutron cancer therapy	2
DO experiment	6

A large collection of requests falling under Computer System Management

References

[1] D. Bogert, Proceedings of the Second International Conference on Accelerator Control Systems, 1985, pp. 8-24.

[2] R. Johnson, Proceedings of the IEEE 1987 Particle Accelerator conference (to be published), and Fermilab Accelerator Division Software Release #159.